SvedbergOpen
DISSEMINATION OF KNOWLEDGE

International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: https://www.svedbergopen.com/

Research Paper Open Access

Building Compliance-Ready Digital Experience Platforms Using Adobe Experience Manager and Artificial Intelligence

Sham Sunder Reddy Dadi

University of Central Missouri, USA. ORCID ID: 0009-0003-8649-1962

Abstract
Regulated enterprises now publish content that carries legal, financial, or clinical weight, but compliance review in most digital experience platforms still happens at the end of the pipeline, as a final check rather than a built-in property of the system. This article proposes a lifecycle governance model for Adobe Experience Manager-based platforms in which metadata design, workflow structure, AI-assisted pre-review, and audit traceability function as architectural decisions rather than downstream inspection steps. The argument draws on enterprise content management theory, information governance and governance-risk-compliance research, and recent empirical work on AI-assisted regulatory review to make the case that distributing compliance controls across the authoring-to-publishing lifecycle reduces both bottlenecks and audit exposure more reliably than a single pre-publish gate. A natural objection follows: AI-assisted pre-review could introduce new risk through false negatives or reviewer over-reliance on automated flags. The article takes this objection seriously and argues that bounded human-AI collaboration, with accountability and traceability built in from the start, offers a more defensible middle path than either fully manual or fully automated review. It closes by identifying what empirical validation of the model would require.

Keywords: digital experience platform governance; Adobe Experience Manager; content compliance; artificial intelligence content review; digital asset management; audit traceability

1. Introduction

A product disclosure page, a customer notice, or a clinical content item published through a digital experience platform is rarely just content. In financial services, insurance, and life sciences, it is a governed artifact, subject to legal review, disclosure obligations, and eventual audit. When one region's page still shows an outdated rate or a lapsed disclaimer while another has been updated, the resulting exposure extends well past a routine publishing mistake.

Most organizations still handle this risk the same way: compliance review sits near the end of the content lifecycle, right before publishing. That arrangement worked reasonably well when publishing volumes were modest and content teams were small enough to know every page personally. At enterprise scale, with thousands of pages and assets moving across business units and regions at once, the same arrangement starts to work against itself. Risk detection gets concentrated at one late point in the process, and the resulting bottleneck pushes reviewers toward speed rather than depth, the opposite of what a compliance gate is supposed to buy an organization.

The pattern is visible across sectors with different regulatory logics. A bank's product-disclosure page answers to financial regulators; a pharmaceutical company's clinical-content page answers to health authorities and to internal medical, legal, and regulatory affairs review; an insurer's policy page answers to state-level insurance regulation. What these settings share is not the specific rule set but the structural fact that content is produced faster than any single end-of-pipeline reviewer can safely absorb it, and that the review step arrives after most of the architectural decisions, about metadata, ownership, and routing, have already been made informally or not at all.

Adobe Experience Manager (AEM) is the platform many of these enterprises already run: author, publish, and dispatcher environments, role-based permissions, replication controls, and a versioned content repository. These are structural primitives, not add-ons, and they are capable of supporting governance earlier in the lifecycle than most implementations actually use them for. Artificial intelligence adds a further piece: content can be screened against approved terminology, required disclosures, and prior compliance patterns before a human reviewer ever opens the file.

The question this article takes up is specific. Can an AEM-based platform, paired with AI-assisted review, be architected so that compliance readiness is continuous and traceable across the lifecycle, rather than parked at one checkpoint near the end? The position argued here is that compliance-readiness belongs in metadata structure, workflow routing, and AI-assisted triage design, not bolted on afterward as an inspection step. That position invites pushback, and it should. A governance practitioner could reasonably argue that spreading compliance logic across the platform multiplies the number of places something can go wrong, or that AI-assisted triage simply trades human review fatigue for a new problem: automation bias. Both objections get addressed directly in what follows, after the model itself is laid out.

2. Literature Review

2.1 Content Governance and Compliance Risk in Enterprise Digital Platforms

Governance has been a persistent sticking point in enterprise content management (ECM) research since the field took shape in the mid-2000s (Tyrvainen et al., 2006). Tyrvainen et al. (2006) described ECM as inherently interdisciplinary, spanning technology, process, and organizational concerns at once, and Alalwan and Weistroffer's (2012) later review sharpened that picture by identifying governance and strategy as two of the four dimensions most closely tied to research maturity in the field, and, tellingly, the two dimensions organizations most often under-resource relative to the technical tooling itself. That imbalance is close to the exact pattern this article is arguing against.

More recent evidence suggests the imbalance persists outside general ECM contexts as well. A systematic review of enterprise content management adoption specifically in anti-corruption and public-sector settings (Alshammari et al., 2025) identified fourteen recurring adoption factors, among them security, top-management support, training, and legal or legislative compliance, validated against practitioner input rather than derived purely from prior literature. Security and legal compliance ranking alongside more conventional adoption factors such as training and top-management support suggests that compliance considerations are not a peripheral concern layered onto ECM adoption after the fact, but one of the core factors organizations already treat as inseparable from the decision to adopt a platform in the first place, even if platform architecture does not always follow through on that recognition.

Tallon, Ramirez, and Short's distinction between information governance and IT governance (Tallon et al., 2013) reframes the underlying problem usefully: information, in their account, has its own lifecycle, ownership, and risk profile, independent of whatever system happens to store it. Kooper, Maes, and Roos Lindgreen (2011) pushed this further, arguing that information governance needs its own decision rights and accountability lines, separate from infrastructure governance entirely. Abraham, Schneider, and vom Brocke's structured review of data governance research (Abraham et al., 2019) arrives at a compatible six-dimension framework, spanning organizational structures, standardization, and monitoring, that treats governance as something designed into a system's operating model rather than appended to it. Vicente and Mira da Silva's conceptual model for integrated governance, risk, and compliance (Vicente & Mira da Silva, 2011) goes a step further still, arguing that treating governance, risk management, and compliance as three separate organizational functions, each with its own tooling and reporting line, is itself a source of inefficiency and blind spots; an integrated GRC model instead treats them as facets of a single decision structure. Applied to a digital experience platform, the combined implication of this literature is that compliance controls should travel with the content artifact itself, through metadata, ownership tags, lifecycle state, rather than sitting off to the side as a review process layered on top, and that the governance, risk, and compliance logic behind those controls should be designed as one coherent structure rather than three uncoordinated ones.

2.2 Adobe Experience Manager and Structured Content Governance

AEM's own documentation talks about governance through templates, workflows, and digital asset management, but academic treatment of what that architecture actually means for compliance is thin. The governance literature offers something more indirect but still useful. Empirical work on COBIT-based IT governance (Alkhaldi et al., 2017) found that organizations applying the framework's structured control objectives reported measurably stronger perceived governance value than organizations without a comparable framework in place, evidence that formal governance frameworks, applied consistently, produce outcomes distinguishable from ad hoc practice rather than functioning as paperwork exercises. Reviews of ISO/IEC 27001 adoption in information-security management (Culot et al., 2021) point toward a related conclusion specific to security controls: certification and structured control frameworks tend to formalize practices that already existed informally, converting tacit knowledge about what counts as adequate security into auditable, repeatable procedure, which is precisely the conversion a compliance-ready content platform needs to perform for governance more broadly.

Metadata quality research (Kumar et al., 2025) makes the point that governance only works when metadata schemas are consistent and well-defined rather than assembled ad hoc, a requirement AEM's asset and content metadata models can support structurally, though nothing forces that support to be used. Case study work on data governance rollouts (Walsh et al., 2025) finds much the same thing from a different angle: clear ownership assignment and defined escalation paths show up repeatedly as success factors, and both map cleanly onto AEM's workflow and permissions model, provided someone configures them that way rather than assuming the platform handles it by default.

Access control research adds another thread. Systematic reviews of access control in regulated settings like electronic health records (Cobrado et al., 2024), together with broader empirical work on access-control systems in live production environments (Parkinson & Khan, 2022), keep arriving at the same finding: role-based access control cuts unauthorized-modification risk, but only when role definitions stay narrow; loosely defined roles and inconsistent enforcement erode the benefit quickly. That finding lands directly on AEM's author-publish-dispatcher separation. The permission boundary between authoring and publishing is a real, designed capability, but like any access-control system, it is only as strong as how carefully it gets configured.

2.3 Artificial Intelligence in Content Review and Classification

Outside the ECM literature, automated review of regulated content has moved quickly, mostly in legal-technology and financial-technology research. El Hamdani et al. (2021) built a hybrid rule-based and machine-learning pipeline for automated GDPR compliance checking, reporting classification performance the authors judged sufficient to support automated flagging of regulation-relevant clauses for human follow-up, rather than autonomous determination. Lore et al. (2023) took a related but more cautious approach, developing an AI framework explicitly designed to support, not replace, human GDPR decisions, the system's output is framed as a recommendation, never a determination. Financial services shows a parallel pattern: Craja et al. (2020) applied deep learning to financial statement fraud detection, which demonstrates that AI-assisted classification of regulated financial documents is technically mature and already in use near compliance functions, even where it hasn't been formally folded into content-management platforms.

The human-AI collaboration literature adds a note of caution to that promise. In Molina and Sundar's (2022) experimental work on AI content moderation, user trust and perceived legitimacy hinged heavily on whether AI involvement was disclosed and whether a human reviewer stayed visibly in the loop; fully automated moderation, even when it got the call right, was trusted less than a hybrid human-AI process. Werder et al. (2022) make a related point from the systems-design side: responsible AI requires explicit data provenance and traceability, not just accurate output, if it's going to hold up under audit. Read together, this body of work suggests AI-assisted content review earns its place in enterprise governance when it functions as a triage layer with visible human accountability, not as a standalone compliance gate. That constraint shapes the framework developed next.

Table 1 summarizes the reference base underlying this literature review by category, to make the evidentiary structure of the argument transparent rather than leaving the balance between foundational theory, formal frameworks, and recent empirical work implicit.

Table 1. Reference Base by Category

Category	Count	Representative Sources
Foundational works	4	Tyrvainen et al. (2006); Alalwan & Weistroffer (2012); Tallon et al. (2013); Kooper et al. (2011)
Theoretical models, frameworks, and standards	4	Abraham et al. (2019); Vicente & Mira da Silva (2011); Alkhaldi et al. (2017); Culot et al. (2021)
Highly cited works and industry-scale studies	2	Alshammari et al. (2025); Abdurrahman et al. (2024)
Recent empirical studies (2020-present)	11	El Hamdani et al. (2021); Lore et al. (2023); Craja et al. (2020); Molina & Sundar (2022); Werder et al. (2022); and 6 others

3. Methodological Approach: A Lifecycle Model for Compliance-Ready Governance

The framework treats compliance readiness as something distributed across five stages of the AEM content lifecycle, not concentrated in any one of them: metadata-driven authoring, AI-assisted pre-review, risk-tiered

workflow routing, role-based publishing, and audit-layer traceability. Each stage is an architectural decision, not an instruction handed to a reviewer, that distinction is the whole departure from gate-based compliance. The five stages draw directly on the integrated governance-risk-compliance logic described by Vicente and Mira da Silva (2011) and on the formalization argument found in COBIT and ISO/IEC 27001 adoption research (Alkhaldi et al., 2017; Culot et al., 2021): governance frameworks earn their value by converting informal practice into structured, repeatable, auditable procedure, and a content platform's architecture is one of the more concrete places that conversion can actually happen.

Table 2 summarizes the recurring governance challenges identified in Section 2 alongside the specific AEM platform control each challenge maps to, and the literature that motivates the mapping. The table is intentionally read left to right, from challenge to control to evidentiary support, rather than as a checklist, since the argument of this article is that these controls only function as a compliance-ready system when they are designed together rather than adopted piecemeal.

Table 2. Compliance Challenges and Corresponding Platform Controls

Governance Challenge	AEM Platform Control	Supporting Literature
Metadata assembled inconsistently or after the fact	Metadata schemas enforced on templates and DAM folders at authoring time	Kumar et al., 2025; Abraham et al., 2019
Compliance review concentrated at a single late checkpoint	AI-assisted pre-review layer producing triage flags before human review	El Hamdani et al., 2021; Lore et al., 2023; Molina & Sundar, 2022
Uniform review depth regardless of actual risk	Risk-tiered workflow routing with escalation rules	Vicente & Mira da Silva, 2011; Alkhaldi et al., 2017
Unreviewed content reaching production	Role-based publishing separated from authoring and review, enforced via replication controls	Cobrado et al., 2024; Parkinson & Khan, 2022
Inability to reconstruct who approved what and when	Continuous audit-layer logging of metadata, AI flags, approvals, and publishing actions	Werder et al., 2022; Regueiro et al., 2021

Metadata-driven authoring.

Content carries structured metadata from the moment it's authored, not from the moment it's reviewed, fields like content owner, business unit, jurisdiction, compliance category, and expiration date, consistent with the metadata quality dimensions the governance literature has already mapped out (Kumar et al., 2025). In practice, this means metadata schemas attached to templates and DAM folders that block an author from finishing a page or asset upload without filling in the fields governance review will need later. That flips the usual order, where metadata gets patched in during review precisely because it was optional during authoring. In regulated document-heavy settings such as pharmaceutical regulatory submissions, this same logic, structuring content and its metadata at the point of creation rather than retrofitting structure later, has already been documented as a driver of efficiency in regulatory data exchange (Beierle et al., 2023), and the parallel to enterprise digital-experience content is direct: structure paid for early is cheaper than structure imposed under deadline pressure.

AI-assisted pre-review.

Once metadata is complete, an AI layer screens the content before a human compliance reviewer sees it, flagging missing disclosures, deviations from approved terminology, and patterns statistically linked to past compliance findings, in line with the hybrid rule-based and machine-learning approaches already demonstrated in regulatory-compliance research (El Hamdani et al., 2021; Lore et al., 2023). The design choice that matters most here, and it follows directly from the trust findings in human-AI moderation research (Molina & Sundar, 2022), is that the AI layer produces a triage recommendation with visible reasoning, which terms or fields tripped the flag, rather than a pass/fail verdict. This is the direct answer to the objection raised in the introduction: AI output functions as an input to human judgment, not a stand-in for it, which narrows the automation-bias risk without pretending to erase it. In practical terms, the pre-review layer combines the two mechanisms most common in the regulatory-text literature: a rule engine that checks for required fields,

prohibited terms, and structural completeness, and a supervised text classifier, of the kind used for clause-level GDPR flagging (El Hamdani et al., 2021), that scores drafted language against prior examples of approved and rejected content. Named-entity recognition over the metadata fields established in Stage 1 lets the classifier condition its output on jurisdiction and content category, so a clause flagged as risky for one region's disclosure requirements is not automatically flagged for a jurisdiction where the same language is unproblematic.

Risk-tiered workflow routing.
Treating every content item to the same review depth is inefficient, and it is also, somewhat counterintuitively, a source of risk in its own right, since reviewer attention is finite and a uniform gate spreads it evenly regardless of where the actual danger sits. AEM's workflow engine routes content according to the risk classification produced during metadata tagging and AI pre-review: a general informational page might need nothing more than business-owner sign-off, while a financial disclosure or clinical item moves through business, legal, and compliance review in sequence, with escalation triggered automatically if a reviewer sits on it past a defined window.

Role-based publishing.
Publishing authority stays separate from authoring and review authority, in keeping with the separation-of-duties principle documented in access-control research (Cobrado et al., 2024; Parkinson & Khan, 2022). AEM enforces this through permission boundaries between author and publish environments and through replication controls that decide what actually reaches the publish tier, regardless of what's been drafted or reviewed upstream. This stage exists for one reason: to stop a single careless or compromised account from pushing unreviewed content into production, exactly the failure mode the access-control literature flags when role definitions run too loose (Parkinson & Khan, 2022).

Audit-layer traceability.
Traceability, in this framework, is not a report generated on request; it's a running property of the system. Every stage above, metadata assignment, AI flags and how they were resolved, workflow approvals, publishing actions, gets logged in enough detail to reconstruct who did what, when, and what the AI recommended versus what the human reviewer actually decided. Work on blockchain-based audit trail design (Regueiro et al., 2021) demonstrates one concrete mechanism for making this kind of log tamper-evident rather than merely append-only, which matters specifically because an audit trail that can be quietly edited after the fact provides little real defensibility, however complete it looks at first glance. Whether a given AEM implementation needs blockchain-grade tamper evidence or a simpler immutable log store is a deployment decision, but the underlying requirement, that the log itself be trustworthy and not just present, holds regardless of mechanism. This directly answers the data-provenance requirement from responsible-AI research (Werder et al., 2022): an AI recommendation that can't be traced back to its inputs and reasoning isn't defensible under regulatory audit, no matter how accurate it turns out to be on average.
Taken as a set, these five stages move compliance from a single decision point into a distributed set of architectural constraints. Each one traces back to a specific requirement surfaced in Section 2's literature, and each maps to an AEM capability that already exists, rather than one this article is proposing be built from scratch. Table 3 restates the five stages as a single reference summary, including the feedback path back into metadata schema design that a mature implementation should eventually build in as compliance findings accumulate.

Table 3. Five-Stage Lifecycle Model Summary

Stage	Architectural Decision	Primary Governance Requirement Addressed
1. Metadata-driven authoring	Mandatory metadata schema enforced at template and DAM level	Consistent, well-defined metadata as a precondition for governance (Kumar et al., 2025)
2. AI-assisted pre-review	Triage flags with visible reasoning, not autonomous verdicts	Defensible AI involvement without displacing human judgment (Molina & Sundar, 2022; Werder et al., 2022)
3. Risk-tiered workflow routing	Review depth scaled to risk classification, with escalation rules	Efficient allocation of finite reviewer attention Vicente & Mira da Silva, 2011; Alkhaldi et al., 2017

Stage	Architectural Decision	Primary Governance Requirement Addressed
4. Role-based publishing	Separation of authoring, review, and publishing permissions	Separation of duties to prevent unreviewed publication Cobrado et al., 2024; Parkinson & Khan, 2022
5. Audit-layer traceability	Continuous, tamper-evident logging across all prior stages	Reconstructable, auditable decision history Werder et al., 2022; Regueiro et al., 2021

A brief illustrative scenario may help ground the model. Consider a life-sciences organization publishing a clinical-trial patient education page through AEM, drawing on content originally structured in a regulated content vault of the kind described in structured pharmaceutical content management research (Beierle et al., 2023). Under a gate-based model, the page would be drafted, routed once to a compliance reviewer near the end of the process, and published or returned. Under the model proposed here, the page is authored against a metadata schema that already requires a jurisdiction tag and an expiration date; an AI layer flags a sentence that deviates from the organization's approved patient-facing terminology library before a human ever sees the page; the flag routes the page into the clinical-and-legal review tier rather than the lighter general-content tier; publishing permission sits with a role distinct from the authoring and reviewing roles; and the entire sequence, including the AI's flag and the reviewer's disposition of it, is logged in a form an auditor could reconstruct months later. None of the individual steps is unusual in isolation. What changes is that each step is a designed platform property rather than something that depends on a diligent reviewer noticing a problem at the last possible moment.

4. Discussion

The claim underneath Section 3 is not obviously true on its face: that spreading compliance logic across the lifecycle beats concentrating it in one thorough pre-publish review. The sharpest counter-argument, raised already in the introduction, is that AI-assisted pre-review creates a new failure mode of its own, false negatives that a tired or under-resourced reviewer might miss precisely because the AI has quietly lowered their guard. That worry has real empirical backing. Automation bias, where people over-trust automated output, shows up consistently in human-AI collaboration research, and Molina and Sundar's (2022) moderation-trust findings indicate that simply knowing AI is involved changes how reviewers calibrate their own judgment, for better or worse.

Table 4 lays out the comparison implicit in this argument more directly, contrasting the gate-based model most organizations still run against the lifecycle-distributed model proposed here, across the dimensions the counter-arguments above actually turn on.

Table 4. Gate-Based Versus Lifecycle-Distributed Compliance Models

Dimension	Gate-Based Model	Lifecycle-Distributed Model
Point of risk detection	Concentrated near publishing	Distributed across authoring, pre-review, routing, and publishing
Reviewer attention allocation	Uniform across all content	Scaled to risk tier
Role of AI (where present)	Often absent or used as a final scan	Triage layer with visible reasoning, positioned before human review
Accountability under failure	Diffuse; hard to reconstruct decision history	Explicit; logged at each stage, human-held by design
Metadata timing	Frequently retrofitted during review	Required at authoring

The framework's answer to that worry is structural rather than a matter of asking reviewers to try harder. Because the AI layer in Section 3 surfaces reasoning instead of verdicts, and because every flag and every resolution gets logged under the audit-traceability stage, the model does not need reviewers to trust the AI correctly in the moment, it needs the organization to be able to tell, after the fact, whether reviewers are rubber-stamping AI output without exercising independent judgment. That doesn't make automation bias disappear,

but it turns an invisible risk into a measurable one, and those are different governance postures entirely. A second objection, related but distinct, concerns speed: risk-tiered routing and multi-stage logging could plausibly slow publishing down relative to one end-of-pipeline check. The framework's response is that most content, once correctly classified by risk tier, needs less review depth than a uniform gate applies by default, the net effect should be redistribution of review effort toward high-risk content, not a blanket increase in review burden across everything. That claim, more than any other in this article, is the one most in need of empirical testing, a point picked up again in Section 6.

Some indirect support for this redistribution claim comes from outside the ECM literature proper. Empirical research on Indonesian banks (Abdurrahman et al., 2024) finds that governance, risk, and compliance capability does not compete with digital-transformation performance gains but strengthens them, banks that treated GRC and digital-transformation capability as complementary, rather than as a compliance tax on transformation, tended to perform better on the outcomes the study measured. That finding is domain-adjacent rather than a direct test of the framework proposed here, but it is at least consistent with the claim that well-structured governance and operational speed are not inherently in tension, which is the assumption the redistribution argument in this article depends on.

There is a further question worth sitting with: where does accountability land when AI-assisted pre-review turns out to be part of the chain leading to a compliance failure? The provenance requirements from Werder et al. (2022) point toward keeping accountability with the human reviewer and the approving business unit, with the AI system's role limited to a documented, auditable recommendation. That division matters for two reasons, regulatory defensibility, obviously, but also for keeping human oversight meaningful rather than symbolic. Without it, the phrase the AI flagged it as low-risk becomes a convenient place to park blame rather than a genuine input into a human decision. The integrated governance-risk-compliance framing from Vicente and Mira da Silva (2011) is useful again here: if governance, risk, and compliance are treated as one coordinated structure rather than three separate reporting lines, accountability for an AI-assisted pre-review decision has a clearer home than it would if, say, the AI tooling sat under IT governance while the compliance consequence of a missed flag sat entirely under a separate legal or compliance function with no visibility into how the tool was configured.

5. Practical Implications

Platform architects should treat governance requirements as inputs to metadata schema and workflow design during AEM implementation itself, not as a problem to solve afterward with review processes bolted onto a platform that was never built to support them. Retrofitting metadata fields once templates and DAM structures are already live is possible, but it costs considerably more than designing them in from day one, a pattern already visible in regulated document-management settings such as pharmaceutical content operations, where structuring content and metadata early in the lifecycle has been linked to faster downstream regulatory data exchange (Beierle et al., 2023) rather than treated as pure administrative overhead.

Compliance and legal teams see a real shift in where their expertise goes under this model. Instead of reviewing every item at the same depth, they're positioned to define the risk-classification rules and terminology libraries that drive AI pre-review and workflow routing in the first place, and to concentrate direct review effort on the tier that actually carries the risk. That requires tighter collaboration between compliance and platform architecture than the traditional gate-based model ever demanded, since under the old model compliance review happened largely independent of how the platform itself was built.

For organizations weighing AI-assisted content review, the practical takeaway is sequencing: disclosure and traceability design should come before automation design, not after. Rolling out an AI pre-review capability without first deciding how its recommendations will be logged, disclosed to reviewers, and made auditable risks recreating exactly the trust and accountability problems the human-AI collaboration literature has already documented (Molina & Sundar, 2022; Werder et al., 2022), regardless of how accurate the underlying model turns out to be in testing.

Data and IT governance teams, finally, have a direct stake in how the metadata schemas underneath Stage 1 get designed, since those schemas are shared infrastructure that both compliance workflows and broader analytics or personalization efforts will eventually draw on. Designing metadata fields narrowly, to satisfy only the immediate compliance use case, tends to produce schemas that later need to be redesigned when a second use case, such as content analytics or audience personalization, needs fields the compliance-only design never anticipated. Coordinating metadata schema design across governance and broader data-strategy stakeholders at the outset avoids that rework.

6. Conclusion

The case made here is that compliance readiness in AEM-based digital experience platforms works better as a distributed architectural property than as a single review checkpoint, and the five-stage model developed above, metadata-driven authoring, AI-assisted pre-review, risk-tiered routing, role-based publishing, audit-layer traceability, is one attempt to operationalize that claim. What the model brings together are two literatures that have mostly grown apart: enterprise content governance research, which has worked out the organizational, metadata, and governance-risk-compliance requirements for defensible governance, and AI-assisted regulatory review research, which has shown technical feasibility but has said comparatively little about how such systems fit into content platforms enterprises are already running. The real contribution is not any single stage of the model, most of them have some precedent in one literature or the other, but the explicit allocation of accountability between AI recommendation and human decision across a full content lifecycle, worked out at the level of platform architecture rather than review policy.

7. Limitations and Future Research

None of this has been tested against operational data from an actual AEM deployment, and the model's central efficiency claim, that risk-tiered routing redistributes review burden rather than adding to it, remains a hypothesis, not a demonstrated result. A natural next step would be case study work comparing review-cycle time and audit-finding rates before and after organizations adopt risk-tiered content workflows. There's a second gap worth naming directly: the evidence behind the AI-assisted pre-review component comes from adjacent domains, GDPR compliance checking, financial fraud detection, online content moderation, rather than from studies of AI pre-review inside enterprise digital-experience-platform workflows specifically. Whether the trust and accountability findings from those adjacent domains carry over cleanly to an AEM-specific implementation is genuinely an open question. And because the framework was built around one vendor's architecture, it's worth asking in future work whether the five-stage model holds up on other enterprise content platforms with different governance primitives built in.

Ethics Statement

Conflict of Interest:

The author declares no conflict of interest.

Statement of Human and Animal Rights:

This study did not involve human participants, animal subjects, or identifiable personal data requiring ethics board review.

Informed Consent:

Not applicable.

References

1. Abdurrahman, A., Gustomo, A., & Prasetio, E. A. (2024). Enhancing banking performance through dynamic digital transformation capabilities and governance, risk management, and compliance: Insights from the Indonesian context. The Electronic Journal of Information Systems in Developing Countries, 90(2), Article e12299. https://doi.org/10.1002/isd2.12299
2. Abraham, R., Schneider, J., & vom Brocke, J. (2019). Data governance: A conceptual framework, structured review, and research agenda. International Journal of Information Management, 49, 424-438. https://doi.org/10.1016/j.ijinfomgt.2019.07.008
3. Alalwan, J. A., & Weistroffer, H. R. (2012). Enterprise content management research: A comprehensive review. Journal of Enterprise Information Management, 25(5), 441-461. https://doi.org/10.1108/17410391211265133
4. Alkhaldi, F. M., Hammami, S. M., & Uddin, M. A. (2017). Understanding value characteristics toward a robust IT governance application in private organizations using COBIT framework. International Journal of Engineering Business Management, 9, 1-8. https://doi.org/10.1177/1847979017703779
5. Alshammari, B. O., Mokhtar, U. A., & Abdulmanap, N. (2025). A systematic literature review on the enterprise content management system adoption to combat corruption. SAGE Open, 15(1), 1-20. https://doi.org/10.1177/21582440251321231
6. Beierle, J., Algorri, M., Cortes, M., Cauchon, N. S., Lennard, A., Kirwan, J. P., Oghamian, S., & Abernathy, M. J. (2023). Structured content and data management-enhancing acceleration in drug development

through efficiency in data exchange. AAPS Open, 9, Article 11. https://doi.org/10.1186/s41120-023-00077-6
7. Cobrado, U. N., Sharief, S., Regahal, N. G., Zepka, E., Mamauag, M. B., & Velasco, L. C. (2024). Access control solutions in electronic health record systems: A systematic review. Informatics in Medicine Unlocked, 49, Article 101552. https://doi.org/10.1016/j.imu.2024.101552
8. Craja, P., Kim, A., & Lessmann, S. (2020). Deep learning for detecting financial statement fraud. Decision Support Systems, 139, Article 113421. https://doi.org/10.1016/j.dss.2020.113421
9. Culot, G., Nassimbeni, G., Podrecca, M., & Sartor, M. (2021). The ISO/IEC 27001 information security management standard: Literature review and theory-based research agenda. The TQM Journal, 33(7), 76-105. https://doi.org/10.1108/TQM-09-2020-0202
10. El Hamdani, R., Mustapha, M., Restrepo Amariles, D., Troussel, A., Meeus, S., & Krasnashchok, K. (2021). A combined rule-based and machine learning approach for automated GDPR compliance checking. In Proceedings of the Eighteenth International Conference on Artificial Intelligence and Law (ICAIL '21) (pp. 40-49). Association for Computing Machinery. https://doi.org/10.1145/3462757.3466081
11. Kooper, M. N., Maes, R., & Roos Lindgreen, E. E. O. (2011). On the governance of information: Introducing a new concept of governance to support the management of information. International Journal of Information Management, 31(3), 195-200. https://doi.org/10.1016/j.ijinfomgt.2010.05.009
12. Kumar, V., Chandrappa, & Harinarayana, N. S. (2025). Exploring dimensions of metadata quality assessment: A scoping review. Journal of Librarianship and Information Science, 57(3), 661-673. https://doi.org/10.1177/09610006241239080
13. Lore, F., Basile, P., Appice, A., de Gemmis, M., Malerba, D., & Semeraro, G. (2023). An AI framework to support decisions on GDPR compliance. Journal of Intelligent Information Systems, 61(2), 541-568. https://doi.org/10.1007/s10844-023-00782-4
14. Molina, M. D., & Sundar, S. S. (2022). When AI moderates online content: Effects of human collaboration and interactive transparency on user trust. Journal of Computer-Mediated Communication, 27(4), Article zmac010. https://doi.org/10.1093/jcmc/zmac010
15. Parkinson, S., & Khan, S. (2022). A survey on empirical security analysis of access-control systems: A real-world perspective. ACM Computing Surveys, 55(6), Article 123. https://doi.org/10.1145/3533703
16. Regueiro, C., Seco, I., Gutierrez-Agueero, I., Urquizu, B., & Mansell, J. (2021). A blockchain-based audit trail mechanism: Design and implementation. Algorithms, 14(12), Article 341. https://doi.org/10.3390/a14120341
17. Tallon, P. P., Ramirez, R. V., & Short, J. E. (2013). The information artifact in IT governance: Toward a theory of information governance. Journal of Management Information Systems, 30(3), 141-178. https://doi.org/10.2753/MIS0742-1222300306
18. Tyrvainen, P., Paivarinta, T., Salminen, A., & Iivari, J. (2006). Characterizing the evolving research on enterprise content management. European Journal of Information Systems, 15(6), 627-634. https://doi.org/10.1057/palgrave.ejis.3000648
19. Vicente, P., & Mira da Silva, M. (2011). A conceptual model for integrated governance, risk and compliance. In H. Mouratidis & C. Rolland (Eds.), Advanced Information Systems Engineering (CAiSE 2011, Lecture Notes in Computer Science, Vol. 6741, pp. 199-213). Springer. https://doi.org/10.1007/978-3-642-21640-4_16
20. Walsh, M. J., McAvoy, J., & Sammon, D. (2025). The data governance journey in practice: Insights from case study research. Information Systems Management, 42(3), 471-488. https://doi.org/10.1080/10580530.2025.2477459
21. Werder, K., Ramesh, B., & Zhang, S. (2022). Establishing data provenance for responsible artificial intelligence systems. ACM Transactions on Management Information Systems, 13(2), Article 22. https://doi.org/10.1145/3503488